arXiv:1611.05517v2 [math.PR] 19 Nov 2016

Lifting linear preferential attachment trees yields the arcsine coalescent

Helmut H. Pitters

Department of Statistics
University of California, Berkeley

November 22, 2016

Abstract

We consider linear preferential attachment trees which are specific scale-free trees also known as (random) plane-oriented recursive trees. Starting with a linear preferential attachment tree of size n we show that repeatedly applying a so-called lifting yields a continuous-time Markov chain on linear preferential attachment trees. Each such tree induces a partition of $\{1,\dots,n\}$ by placing labels in the same block if and only if they are attached to the same node in the tree. Our main result is that this Markov chain on linear preferential attachment trees induces a partition valued process which is equal in distribution (up to a random time-change) to the arcsine n-coalescent, that is the multiple merger coalescent whose Λ measure is the arcsine distribution.

1 Introduction

A linear preferential attachment tree T of size n is a random planar rooted tree on n nodes labeled $1,\dots,n$ such that the labels along any non-backtracking path starting from the root are increasing. The tree T can be constructed by attaching nodes as they arrive in the order of increasing labels as follows.

1. Start with a root node labeled 1.
2. At step $n-1$ we have a (random) tree on $n-1$ nodes with labels $1,\dots,n-1$. A new node with label n is added to the existing tree, namely it is attached with an edge to node v with probability proportional to

$$d^+(v)+1,$$

where $d^+(v)$ counts the number of successors of v. There are $d^+(v)+1$ available positions for an additional successor of v, since T is planar. Node n is assigned to one of these positions chosen uniformly at random.

Let us now introduce a continuous-time Markov chain on linear preferential attachment trees. The Markov cain starts in state T and its absorbing state is the tree consisting of only one node. If T has $n > 1$ nodes, after waiting an exponential time of rate n, among all nodes of T choose one node U uniformly at random. If U is a leaf, do nothing. Otherwise, select one of Us successors, V say, uniformly at random. We now lift the edge $\{U, V\}$ as follows. Collect the labels attached to vertices in the subtree T_V rooted at V and attach them to U, then remove the edge $\{U, V\}$ together with T_V. We call this procedure "lifting" of an edge, following Berestycki's lecture notes [4]. However, in the literature variants of this procedure are sometimes called "cutting" or "pruning". The tree T^{L} obtained after lifting T is again a linear preferential attachment tree, as we show in Lemma 1. Moreover, T^{L} is labeled by the blocks of some partition π of $[n]$. We say that π is induced by T^{L}. A moment's thought shows that, in our context, the appropriate way to order the blocks of this partition is according to their least element. Applying this lifting procedure repeatedly yields a Markov chain taking values in the set of linear preferential attachment trees which are labeled by the blocks of partitions of $[n]$. If we start with the linear preferential attachment tree T of size n and only keep track of the partitions of $[n]$ induced by the lifted trees, we obtain a stochastic process with state space the set $\mathscr{P}_{[n]}$ of partitions of $[n]$ such that the partitions become coarser and coarser as time passes. In fact, starting with a linear preferential attachment tree of size n our main result, Theorem 1, shows that by repeated lifting we obtain (up to a random time change) the arcsine n-coalescent, which is the multiple merger coalescent corresponding to the measure Λ taken to be the arcsine distribution. See section 2 for a definition of multiple merger coalescent processes.

The first construction of a coalescent process by a similar lifting procedure applied to random recursive trees was given by Goldschmidt and Martin [6]. The authors start with a random recursive tree and show that the partition-valued process induced by repeated lifting yields the Bolthausen-Sznitman n-coalescent corresponding to the Λ measure given by the uniform distribution. Abraham and Delmas give another construction of the beta$(\frac{3}{2}, \frac{1}{2})$ n-coalescent by lifting random binary trees in [1], and a construction of the jump chain of the beta$(1+\alpha, 1-\alpha)$ n-coalescent by lifting stable Galton-Watson trees in [2]. It should however be noted that the lifting procedures employed in these examples differ from each other.

2 Main Results

An *increasing tree* on the labels $1, \dots, n$ is a rooted tree on n nodes which are labeled by $1, \dots, n$ such that any sequence of labels along any non-backtracking path starting at the root is increasing. A *plane-oriented recursive tree (PORT)* is a planar increasing tree, i.e. the successors of any node are ordered. A PORT on n nodes can be constructed recursively as follows.

1. Start with the tree t_1 consisting only of the root with label 1, which trivially is a PORT.

2. Given a PORT t_{n-1} on $n-1$ nodes pick a node v in t_{n-1} and put a further node with label n into any of the $d^+(v)+1$ positions available at v.

Denote by $\#A$ the cardinality of a set A. We slightly abuse notation and write $v \in t$ if v is a node of t. All $n-1$ nodes of t_{n-1} except for the root are successors of some node, thus there are

$$\sum_{v \in t_{n-1}} (d^+(v)+1) = n-2+n-1 = 2n-3 = 2(n-1)-1$$

PORTs on n nodes that can be constructed by adding node n to t_{n-1}. If t_n, t'_n are PORTs of size n constructed in this way from PORTs t_{n-1}, respectively t'_{n-1}, of size $n-1$, then $t_{n-1} \neq t'_{n-1}$ clearly implies $t_n \neq t'_n$. Consequently, letting $\mathfrak{P}_n$ denote the *set of plane-oriented recursive trees on n labeled nodes*, we just showed that $\#\mathfrak{P}_n = (2(n-1)-1)\#\mathfrak{P}_{n-1}$, hence

$$\#\mathfrak{P}_n = 1 \cdot 3 \cdots (2(n-1)-1) = (2(n-1)-1)!! = \frac{1}{2^{n-1}} \frac{(2(n-1))!}{(n-1)!} = \frac{n!}{2^{n-1}} C_{n-1}, \tag{1}$$

where for any integer $n \geq -1$ the *double factorial* is defined by

$$n!! \coloneqq \begin{cases} 1 \cdot 3 \cdot 5 \cdots (n-2)n & \text{if } n \text{ is odd,} \\ 2 \cdot 4 \cdot 6 \cdots (n-2)n & \text{if } n \text{ is even,} \\ 1 & \text{if } n \in \{-1, 0\}, \end{cases} \tag{2}$$

and $C_n \coloneqq (2n)!/(n!(n+1)!)$, $n \in \mathbb{N}_0$, denotes the *nth Catalan number*.

A *linear preferential attachment tree (LPAT) of size n* is an element T_n of $\mathfrak{P}_n$ drawn uniformly at random. At times we write LPAT(n) for "LPAT of size n," respectively PORT(n) for "PORT of size n". Notice that the recursive construction of a PORT(n) immediately yields a recursive construction of an LPAT(n) T_n by picking a node v in an LPAT($n-1$) with probability proportional to $d^+(v)+1$, and attaching a new node with label n to v.

Remark 1. (i) Fix two natural numbers m, n with $m \leq n$. Define the map ρ_{nm} from $\mathfrak{P}_n$ to $\mathfrak{P}_m$, which we call *restriction*, as follows. If $t_n \in \mathfrak{P}_n$ is a plane-oriented recursive tree of size n, let $\rho_{nm}(t_n)$ be the subtree in t_n spanned by the nodes whose labels are smaller than or equal to m. If T_n is an LPAT(n) and T_m is an LPAT(m), it follows from the recursive construction of linear preferential attachment trees that

$$\rho_{nm}(T_n) =_d T_m. \tag{3}$$

(ii) Consider a sequence $\{T_k, 1 \leq k \leq n\}$, where T_k is an LPAT of size k, and T_{k+1} is obtained from T_k by the recursive (random) construction we just described. Clearly, $\{T_k\}$ is a Markov chain. Moreover, T_k contains all the information about the past $\{T_1, \ldots, T_k\}$, since obviously T_j, $j \leq k$, is the subtree in T_k spanned by the nodes with labels $1, \ldots, j$, and T_k is therefore a representation of the σ-algebra generated by the $T_1, \ldots, T_k$.

A *partition* of a nonempty set A is a set, π say, of nonempty pairwise disjoint subsets of A whose union is A. The members of π are called the *blocks* of π. Let $\mathscr{P}_A$ denote the set of all partitions of A.

In what follows we want somewhat more flexibility in the labeling of trees. Namely, we want to label the nodes in a tree by blocks B of a partition of $[n]$. To this end we endow any partition π of $[n]$ by the *order of least elements*, denoted $\leq$, namely let $B \leq C$ if and only if $\min B \leq \min C$ for any two blocks B, C in π. If t is a tree whose nodes are labeled by the blocks in π we call $\mathfrak{L}(t) \coloneqq \pi$ the *label set of* t. With this definition a *plane-oriented recursive tree on* $\pi = \{B_1, \ldots, B_k\} \in \mathscr{P}_n$, PORT($\pi$) for short, respectively a *linear preferential attachment tree on* π, LPAT(π) for short, is a plane-oriented recursive tree, respectively linear preferential attachment tree on k nodes which are labeled by the blocks $B_1, \ldots, B_k$. Denote by $\mathfrak{P}_\pi$ the *set of all plane-oriented recursive trees on* $\#\pi$ *nodes labeled by the blocks* $B_1, \ldots, B_k$.

Remark 2. (i) Fix two natural numbers m, n such that $m \leq n$. Fix a partition $\pi \in \mathscr{P}_n$, and let π' be the restriction of π to $[m]$. Moreover, define the map $\rho_{\pi m} \colon \mathfrak{P}_\pi \to \mathfrak{P}'_\pi$ as follows. If T is an LPAT(π), consider the subtree $\bar{T}$ of T spanned by the nodes whose labels (which are subsets of $[n]$) contain an element in $[m]$. Moreover, let T' be obtained from $\bar{T}$ by restricting the labels in $\bar{T}$ to $[m]$, i.e. by replacing any label B of a node $v \in \bar{T}$ by $B \cap [m]$. We set

$$\rho_{\pi m}(T) \coloneqq T', \tag{4}$$

and call T' the *restriction of* T *to* m. Notice that, again from the recursive construction of linear preferential attachment trees, for any LPAT(π) T and any LPAT(π') T' we have

$$\rho_{\pi m}(T) =_d T'. \tag{5}$$

(ii) Notice that in our nomenclature a linear preferential attachment tree of size n is a tree chosen uniformly at random among all plane-oriented recursive trees of size n. In particular, a PORT is always a deterministic tree, whereas for any $n \geq 3$ an LPAT(n) is a random tree. However, in the literature the reader may occasionally find that the term PORT is used for both, deterministic plane-oriented recursive trees as we define them, as well as for linear preferential attachment trees. Other names for LPATs that appear in the literature are heap-ordered trees, nonuniform recursive trees and scale-free trees, cf. [7].

We now define the operation of lifting which is at the heart of our construction of the arcsine coalescent. Consider a rooted tree t on n labeled nodes. *Lifting an edge $e = \{u, v\}$ in t* works as follows: Assume that u is closer (in graph distance) to the root than v. Then attach the labels on the subtree t_v rooted at v to u, discard both t_v and e, and only keep track of the subtree containing the root of t. In what follows we will choose the edge that is to be lifted in a particular and random fashion. Namely, we pick a node U in t uniformly at random, and, provided U is not a leaf, we pick one of Us successors, call it V, uniformly at random. By *lifting a tree t* we mean picking an edge $\{U, V\}$ randomly in the manner described above and than lifting $\{U, V\}$ in t. We denote by t^{L} the tree that is obtained by lifting t.

Our first observation is that if we lift $\{U, V\}$ in an LPAT T_n of size n we obtain an LPAT T_n^{L} on the new label set $\mathfrak{L}(T_n^{\mathsf{L}})$.

Lemma 1. *Let T_n be an LPAT of size n and let T_n^{L} be the tree obtained by lifting T_n. Then, conditionally on T_n^{L} having label set π, T_n^{L} is an LPAT(π).*

Proof. Fix a subset $C \subseteq [n]$ of size $k := \#C \geq 2$. Let t_C denote an arbitrary but fixed plane-oriented recursive tree of size $n - k + 1$ with label set $\langle C; n\rangle$, which is the partition of $[n]$ whose blocks consist of the elements in $[n] \setminus C$ and the block C. Then

$$\begin{aligned}\mathbb{P}\left\{T_n^{\mathsf{L}} = t_C \mid T_n^{\mathsf{L}} \text{ has label set } \langle C; n\rangle\right\} &= \frac{\mathbb{P}\left\{T_n^{\mathsf{L}} = t_C\right\}}{\mathbb{P}\left\{T_n^{\mathsf{L}} \text{ has label set } \langle C; n\rangle\right\}} \\ &= \#\mathfrak{P}_{n-k+1}^{-1},\end{aligned}$$

where the last equality is seen as follows. Slightly abusing notation, we write $d^+_{t_C}(C)$ for the number of successors of the node in t_C that carries label C. By the recursive construction there are $(d^+_{t_C}(C)+1)\#\mathfrak{P}_{k-1}$ PORTs of size n whose subtree consisting of the first $n-k+1$ nodes agrees with t_C (where we identify the labels $c := \min C$ and C), each of which is equally likely to be observed under T_n, i.e.

$$\begin{aligned}\mathbb{P}\left\{T_n^{\mathsf{L}} = t_C\right\} &= \mathbb{P}\left\{T_n^{\mathsf{L}} = t_C | T_n \text{ contains } t_C\right\} \mathbb{P}\left\{T_n \text{ contains } t_C\right\} \qquad (6)\\ &= \frac{1}{n(d^+_{t_C}(C)+1)} \frac{(d^+_{t_C}(C)+1)\#\mathfrak{P}_{k-1}}{\#\mathfrak{P}_n} = \frac{1}{n}\frac{\#\mathfrak{P}_{k-1}}{\#\mathfrak{P}_n}.\end{aligned}$$

To lift the one edge in T_n which connects the nodes with labels c and $c' \coloneqq C \setminus \{c\}$ (c' is the root of the subtree with label set $C \setminus \{c\}$) provided T_n contains t_C, one first has to pick the node with label c, which happens with probability $1/n$ and, conditionally on c being picked, one has to pick the successor of c with label c', which happens with probability $1/(d^+_{t_C}(C)+1)$. Consequently, the probability to pick the one edge in T_n that yields t_C after being lifted is $1/(n(d^+_{t_C}(C)+1))$. It is important to notice that the probability in (6) only depends on C via $\#C$.

Moreover,

$$\begin{aligned}\mathbb{P}\left\{T_n^{\mathsf{L}} \text{ has label set } \langle C; n\rangle\right\} &= \sum_{t'_C \in \mathfrak{P}_{\langle C;n\rangle}} \mathbb{P}\left\{T_n^{\mathsf{L}} = t'_C\right\} \\ &= \frac{1}{n}\frac{\#\mathfrak{P}_{n-k+1}\#\mathfrak{P}_{k-1}}{\#\mathfrak{P}_n}.\end{aligned}$$

The claim follows. □

To summarize, starting with an LPAT of size n repeated lifting yields a Markov chain on $\bigcup_{\pi\in\mathscr{P}_n}\mathcal{L}_\pi$, where $\mathcal{L}_\pi$ denotes the set of linear preferential attachment trees on π. We now modify this lifting process to obtain a continuous-time Markov chain $\{\mathscr{T}_n(t), t \geq 0\}$ as follows. Start with initial state $\mathscr{T}_n(0) = T_n$, where T_n is an LPAT of size n. If $\mathscr{T}_n$ is in state T, attach an exponential clock to each node v of T that rings at rate 1, all clocks being independent and independent of T. When the first clock rings, if the corresponding node V is a leaf, do nothing. Otherwise, pick a successor U of V uniformly at random and lift the edge $\{U, V\}$ in T to obtain the next state T^{L}.

Before we state our main result, we recall the notion of multiple merger coalescent processes which were introduced independently by Donnelly and Kurtz [5], Pitman [9] and Sagitov [11]. A *(standard) multiple merger n-coalescent* Π_n is a continuous-time Markov chain with state space $\mathscr{P}_{[n]}$, the set of all partitions of $[n] \coloneqq \{1,\ldots,n\}$, and initial state $\Delta_n \coloneqq \{\{1\},\{2\},\ldots,\{n\}\}$ such that if Π_n is in a state of b blocks any $1 \leq k \leq b$ specific blocks merge at rate

$$\lambda_{b,k} \coloneqq \int_0^1 x^{k-2}(1-x)^{b-k}\Lambda(dx), \tag{7}$$

where Λ is a finite measure on the unit interval. The process Π_n is also referred to as the Λ n-coalescent. The integral formula (7) for the transition rates is due to Pitman [9] and follows from the requirement that the Λ n-coalescents be consistent as n varies. Here consistency means that for each $m \leq n$ the restriction of Π_n to $[m]$ is equal in distribution to Π_m. In particular, there exists a process $\Pi \coloneqq \{\Pi(t), t \geq 0\}$, the so-called Λ *coalescent*, with state space the partitions of the positive integers $\mathbb{N} \coloneqq \{1, 2, \ldots\}$ such that for any $n \in \mathbb{N}$ the restriction of Π to $[n]$ is equal in distribution to Π_n.

For $a, b > 0$ the beta(a, b) coalescent is the multiple merger coalescent corresponding to Λ the *beta distribution with parameters a, b* on $(0, 1)$ with density

$$x \mapsto \frac{x^{a-1}(1-x)^{b-1}}{B(a,b)}\mathbf{1}_{(0,1)}(x) \qquad (x \in \mathbb{R}) \tag{8}$$

where $B(a, b) := \int_0^1 x^{a-1}(1-x)^{b-1}dx = \Gamma(a)\Gamma(b)/\Gamma(a+b)$ denotes the *beta integral* and for $x > 0$ the *gamma function* is defined by $\Gamma(x) := \int_0^\infty t^{x-1}e^{-t}dt$. Denoting by $\lambda_{n,k}(a, b)$ the infinitesimal rates of the beta(a, b) coalescent, equation (7) implies

$$\lambda_{n,k}(a,b) = \frac{B(n-k+b, k-2+a)}{B(a,b)}. \tag{9}$$

Since the beta distribution with parameters $\frac{1}{2}, \frac{1}{2}$ and density

$$x \mapsto \frac{1}{\pi\sqrt{x(1-x)}}\mathbf{1}_{(0,1)}(x) \qquad (x \in \mathbb{R}) \tag{10}$$

is the *arcsine law*, we call the corresponding multiple merger coalescent the *arcsine coalescent*. By $\mathbf{1}_A(x)$ we denote the indicator of a set A which equals one if $x \in A$ and zero otherwise. Using $\Gamma(n+\frac{1}{2}) = \sqrt{\pi}(2n)!/(4^n n!)$ and (7) we find the transition rates of the arcsine coalescent to be given by

$$\begin{aligned}\lambda_{n,k}(\frac{1}{2},\frac{1}{2}) &= \frac{B(k-2+\frac{1}{2}, n-k+\frac{1}{2})}{B(\frac{1}{2},\frac{1}{2})} = \frac{1}{\pi}\frac{\Gamma(k-2+\frac{1}{2})\Gamma(n-k+\frac{1}{2})}{\Gamma(n-1)} \\ &= \frac{(2(k-2))!}{4^{k-2}(k-2)!}\frac{(2(n-k))!}{4^{n-k}(n-k)!} = \frac{(k-1)!(n-k+1)!}{4^{n-2}}C_{k-2}C_{n-k}\end{aligned} \tag{11}$$

for $n \geq k \geq 2$.

We now turn to the stochastic process Π'_n recording the partitions of $[n]$ induced by the process $\mathscr{T}_n$ of repeatedly lifting linear preferential attachment trees. More formally, $\Pi'_n := \{\Pi'_n(t), t \geq 0\}$ is defined by letting $\Pi'_n(t) := \mathfrak{L}(\mathscr{T}_n(t))$. The next theorem shows that up to a random time change the process Π'_n is the arcsine coalescent restricted to $[n]$.

Theorem 1. *The process $\Pi'_n := \{\Pi'_n(t), t \geq 0\}$ defined by $\Pi'_n(t) := \mathfrak{L}(\mathscr{T}_n(t))$ is a continuous-time Markov chain with state space the partitions of $[n]$, initial state Δ_n and absorbing state $\{[n]\}$ such that whenever Π'_n is in a state of b blocks a merger of k blocks occurs at rate*

$$\lambda'_{b,k} = \frac{2^{b-2}(b-2)!}{\#\mathfrak{P}_b}\lambda_{b,k}(\frac{1}{2},\frac{1}{2}) = \frac{1}{2(b-1)bC_{b-1}}\lambda_{b,k}(\frac{1}{2},\frac{1}{2}) \qquad (b \geq k \geq 2). \tag{12}$$

Proof. Because of Lemma 1 it suffices to compute the transition rates of Π'_n in its initial state Δ_n. Fix a subset $C \subseteq [n]$ of size $k \coloneqq \#C \geq 2$. Let $\langle C; n\rangle$ denote the partition of $[n]$ consisting of the singletons in $[n] \setminus C$ and the block C. Recall that $\mathscr{T}_n$ starts in state $\mathscr{T}_n(0) = T_n$, where T_n is an LPAT(n). The rate at which we see a lifted tree T_n^{L} whose label set consists of C and the elements in $[n] \setminus C$ is given by

$$\lambda'_{n,k}(C) = \frac{\#\mathfrak{P}_{n-k+1}(d^+_{t_C}(C)+1)\#\mathfrak{P}_{k-1}}{\#\mathfrak{P}_n}\frac{1}{n(d^+_{t_C}(C)+1)}n = \frac{\#\mathfrak{P}_{n-k+1}\#\mathfrak{P}_{k-1}}{\#\mathfrak{P}_n}. \quad (13)$$

In particular, $\lambda'_{n,k}$ only depends on C via $\#C = k$. For this reason we drop the argument C and write $\lambda'_{n,k}$. The first equality in (13) holds since there are $\#\mathfrak{P}_{n-k+1}$ PORTs on $\langle C; n\rangle$, and $\#\mathfrak{P}_{k-1}$ PORTs on the partition of $C \setminus \{c\}$ into singletons, where $c \coloneqq \min C$. Moreover, any PORT(n) built by choosing an element from each of these sets of trees and joining their nodes labeled C, respectively c, by an edge in $d^+_{t_C}(C)+1$ different ways only yields $T_n^{\mathsf{L}} = t_C$ if we lift this particular edge, which happens with probability $1/(n(d^+_{t_C}(C)+1))$. Finally, the overall rate at which we see an event happen when Π'_n is in a state consisting of n blocks is n, hence the last factor.

Recall (cf. [3, Exercise 13.1.14, p. 609]) that the double factorial with odd arguments can be expressed in terms of the gamma function as

$$\Gamma(n+\frac{1}{2}) = \sqrt{\pi}\frac{(2n-1)!!}{2^n} \qquad (n \in \mathbb{N}_0). \quad (14)$$

Use (1) and (14) to rewrite $\lambda'_{n,k}$ as

$$\begin{aligned}
\lambda'_{n,k} &= \frac{\#\mathfrak{P}_{n-k+1}\#\mathfrak{P}_{k-1}}{\#\mathfrak{P}_n} \quad (15)\\
&= \frac{(2(n-k)-1)!!(2(k-2)-1)!!}{\#\mathfrak{P}_n}\\
&= \Gamma(n-k+\frac{1}{2})\frac{2^{n-k}}{\sqrt{\pi}}\Gamma(k-\frac{3}{2})\frac{2^{k-2}}{\sqrt{\pi}}\frac{1}{\#\mathfrak{P}_n}\\
&= \frac{2^{n-2}}{\pi}\frac{(n-2)!}{\#\mathfrak{P}_n}\frac{\Gamma(n-k+\frac{1}{2})\Gamma(k-\frac{3}{2})}{\Gamma(n-1)}\\
&= \frac{2^{n-2}}{\pi}\frac{(n-2)!}{\#\mathfrak{P}_n}B(k-\frac{3}{2}, n-k+\frac{1}{2}),
\end{aligned}$$

and, recalling the transition rates of the arcsine coalescent in (11),

$$\begin{aligned}\lambda'_{n,k} &= \frac{2^{n-2}}{\pi}\frac{(n-2)!}{\#\mathfrak{P}_n}B(\frac{1}{2},\frac{1}{2})\lambda_{n,k}(\frac{1}{2},\frac{1}{2}) \\ &= \frac{2^{n-2}(n-2)!}{\#\mathfrak{P}_n}\lambda_{n,k}(\frac{1}{2},\frac{1}{2}) = \frac{(n-1)!(n-2)!}{2(2(n-1))!}\lambda_{n,k}(\frac{1}{2},\frac{1}{2}).\end{aligned} \tag{16}$$

□

We now turn to the process recording the limiting frequency as $n \to \infty$ of the block in Π_n that contains 1.

3 The block containing 1

In order to better understand the behaviour of the block in Π'_n that contains the label 1 we recall the notion of an exchangeable partition and the Chinese Restaurant Process. A relabeling of a partition π of $[n]$ according to some permutation σ of $[n]$ is the partition $\sigma\pi$ consisting of the blocks

$$\sigma B := \{\sigma(b)\colon b \in B\} \qquad (B \in \pi).$$

A random partition $\bar{\Pi}$ of $[n]$ is called *exchangeable* if its distribution is invariant under any relabeling, i.e. if for any permutation σ of $[n]$ one has

$$\sigma\bar{\Pi} =_d \bar{\Pi}.$$

The partitions we have encountered so far that are induced by lifting LPATs are clearly exchangeable. There are various constructions of exchangeable partitions, and the construction via the Chinese Restaurant Process will turn out to be useful for our purposes. To review the Chinese Restaurant Process we closely follow Pitman's lecture notes [10].

The *Chinese Restaurant Process*, first introduced in Pitman and Dubins cite, is a discrete-time Markov chain whose state at time m is a random permutation σ_m of $[m]$. The cycles of σ_m constitute a (random) partition $\bar{\Pi}_m$ of $[m]$, and these random partitions are *consistent* as m varies, that is for each m and $l \leq m$ the restriction $\bar{\Pi}_m|_{[l]}$ of $\bar{\Pi}_m$ to $[l]$ is equal in distribution to $\bar{\Pi}_l$. Here the restriction $\pi|_B$ of any partition π of A to a subset $B \subseteq A$ is the partition of B consisting of the non-empty blocks $C \cap B$ where C ranges over all blocks in π. Here we focus on a special case of the Chinese Restaurant Process parameterized by a pair of real numbers (α, θ) such that $0 \leq \alpha \leq 1$ and $\theta > -\alpha$. Picture then a restaurant with an unlimited number of empty tables numbered $1, 2, \ldots$, each capable of seating an unlimited number of customers. Customers arrive one by one, they are numbered

in the order of arrival, and take seats according to the following rule: customer 1 sits at table 1. Suppose m customers already arrived and together occupy the first $k \in \mathbb{N}$ tables with $m_i \geq 1$ customers sitting at table $1 \leq i \leq k$. The next customer $m+1$ chooses to sit

1. at table i with probability $(m_i - \alpha)/(m+\theta)$, where he chooses his left neighbor among the customers at table i uniformly at random,

2. alone at the $(k+1)$th table with probability $(\theta + k\alpha)/(m+\theta)$.

If we assign the integers in $1, \ldots, m$ to the same block according to whether or not the corresponding customers sit at the same table, we obtain a random partition of $[m]$. Denote this partition by $\bar{\Pi}_m$. By Kingman's theory of exchangeable partitions there exists a partition $\bar{\Pi}$ of the positive integers $\mathbb{N}$ such that for any $m \in \mathbb{N}$ the restriction of $\bar{\Pi}$ to $[n]$ is equal in distribution to $\bar{\Pi}_m$ and for each block $B \in \bar{\Pi}$ its *asymptotic frequency*

$$\lim_{m\to\infty} \frac{\#(B \cap [m])}{m} \tag{17}$$

exists almost surely. Moreover, it is known that these limiting frequencies in size-biased order of least elements have the representation

$$(\tilde{P}_1, \tilde{P}_2, \ldots) =_d (W_1, \bar{W}_1 W_2, \bar{W}_1 \bar{W}_2 W_3, \ldots), \tag{18}$$

where the $(W_i)_{i\geq 1}$ are independent, W_i is governed by a beta$(1-\alpha, \theta+i\alpha)$ distribution, and $\bar{W}_i := 1 - W_i$. The distribution of $(\tilde{P}_1, \tilde{P}_2, \ldots)$ is the so-called *Griffiths-Engen-McCloskey distribution* with parameters (α, θ), denoted GEM(α, θ). The distribution of $(P_1, P_2, \ldots)$, defined by ranking the $\tilde{P}_1, \tilde{P}_2, \ldots$ in decreasing order, is the so-called *Poisson-Dirichlet distribution* with parameters (α, θ).

There is a Chinese Restaurant Process sitting in an LPAT(n) that we now turn to. We now study another partition $\pi(t)$ (not to be confused with the label set $\mathfrak{L}(t)$ of t) of $[n]$ induced by a tree t with nodes labeled by a partition π of $[n]$. Let ρ denote the root of t, and for any node v in t let t_v denote the subtree in t rooted at v. Define two labels i, $j \in [n]$ to be in the same block of $\pi(t)$ if and only if one of the subtrees of t rooted at a successor of ρ contains both i and j, more precisely, if $i, j \in \{m\colon m \in B \in \mathfrak{L}(t_v)\}$ for some successor v of ρ.

From the recursive construction of LPATs it is immediate that if T_{n+1} is an LPAT of size $n+1$ and T_n is an LPAT of size n, the restriction of $\pi(T_{n+1})$ to $[n]$ is equal in distribution to $\pi(T_n)$ (i.e. the partitions $\pi(T_{n+1})$ and $\pi(T_n)$ are said to be consistent).

Suppose then that T_n is an LPAT of size n with label set $\Delta_{[n]}$ such that its root has k successors each of which subtends a subtree of size $n_i \geq 1$, in particular

$\sum_i n_i = n$. In order to obtain T_{n+1}, we attach the label $n+1$ to a node v in T_n with probability proportional to $d^+(v)+1$. In terms of partitions, the new node $n+1$

1. creates a new block w.p. $\frac{k+1}{2n+1} = \frac{k\frac{1}{2}+\frac{1}{2}}{n+\frac{1}{2}}$,

2. joins a block of size n_i w.p. $\frac{2n_i-1}{2n+1} = \frac{n_i-\frac{1}{2}}{n+\frac{1}{2}}$.

This shows that $\pi(T_n)$ has the same distribution as the partition of $[n]$ induced by the $(\frac{1}{2},\frac{1}{2})$-Chinese Restaurant Process. A more general result revealing the Chinese Restaurant Process in a preferential attachment tree is given in [8, Proposition 3] (notice however that Kuba and Panholzer call these trees generalized plane-oriented recursive trees). The basic properties of the Chinese Restaurant Process reviewed earlier imply that as $n \to \infty$, the asymptotic frequencies $(\tilde{P}_1, \tilde{P}_2, \ldots)$ of the descendencies of the successors of the root (in order of their appearance) follow a GEM$(\frac{1}{2},\frac{1}{2})$ distribution.

Consider now the asymptotic frequency

$$F'(t) := \lim_{n\to\infty} \frac{\#\{1 \le i \le n\colon i \text{ and } 1 \text{ are in the same block of } \Pi'(t)\}}{n} \tag{19}$$

of the block in $\Pi'_n(t)$ containing 1. The magnitudes of the jumps of $\{F'(t), t \ge 0\}$ are given by $(\tilde{P}_i)$, and the interarrival times between successive jumps form an i.i.d. sequence of exponential 1 random variables.

Acknowledgements. The author thanks Steve Evans, Christina Goldschmidt and Martin Möhle for comments on an earlier version of this manuscript, as well as James Martin and Jim Pitman for fruitful discussions.

References

[1] Romain Abraham and Jean-François Delmas, *A construction of a β-coalescent via the pruning of binary trees*, J. Appl. Probab. **50** (2013), no. 3, 772–790. MR 3102514

[2] ———, *β-coalescents and stable Galton-Watson trees*, ALEA Lat. Am. J. Probab. Math. Stat. **12** (2015), no. 1, 451–476. MR 3368966

[3] George B. Arfken and Hans J. Weber, *Mathematical methods for physicists*, fifth ed., Harcourt/Academic Press, Burlington, MA, 2001. MR 1810939

[4] Nathanaël Berestycki, *Recent progress in coalescent theory*, Ensaios Matemáticos [Mathematical Surveys], vol. 16, Sociedade Brasileira de Matemática, Rio de Janeiro, 2009. MR 2574323 (2011d:60225)

[5] P. Donnelly and T. G. Kurtz, *Particle representations for measure-valued population models*, The Annals of Probability **27** (1999), no. 1, 166–205.

[6] C. Goldschmidt and J. Martin, *Random recursive trees and the Bolthausen-Sznitman coalesent*, Electron. J. Probab. **10** (2005), no. 21, 718–745.

[7] Hsien-Kuei Hwang, *Profiles of random trees: plane-oriented recursive trees*, Random Structures Algorithms **30** (2007), no. 3, 380–413. MR 2309623

[8] M. Kuba and A. Panholzer, *On moment sequences and mixed Poisson distributions*, ArXiv e-prints (2014).

[9] J. Pitman, *Coalescents with multiple collisions*, Ann. Probab. **27** (1999), no. 4, 1870–1902. MR 1742892 (2001h:60016)

[10] ———, *Combinatorial stochastic processes*, Springer-Verlag, Berlin, 2006.

[11] S. Sagitov, *The general coalescent with asynchronous mergers of ancestral lines*, J. Appl. Probab. **36** (1999), no. 4, 1116–1125. MR 1742154 (2001f:92019)